DOCUMENT RESUME

ED 400 277 TM 025 532

AUTHOR Parkes, Jay
TITLE Optimal Designs for Performance Assessments: The Subject Factor.
PUB DATE Apr 96
NOTE 22p.; Paper presented at the Annual Meeting of the National Council on Measurement in Education (New York, NY, April 9-11, 1996).
PUB TYPE Reports - Evaluative/Feasibility (142) -- Speeches/Conference Papers (150)

EDRS PRICE MF01/PC01 Plus Postage.
DESCRIPTORS *Cost Estimates; *Performance Based Assessment; *Portfolio Assessment; Portfolios (Background Materials); Program Development; *Sample Size; *Test Construction; Test Reliability; Writing Evaluation
IDENTIFIERS *Large Scale Assessment

ABSTRACT

Much speculation abounds concerning how expensive performance assessments are or are going to be. Recent projections indicate that, in order to achieve an acceptably high generalizability coefficient, many additional tasks may need to be added, which will enlarge costs. Such projections are, to some degree, correct, and to some degree simplistic. The current investigation uses two synthetic examples, based on published costs and variance components, and a constrained optimization procedure to examine the complex relationships among reliability, cost, and sample size. The first example is a limited writing sample situation, and the second is a large-scale portfolio assessment. Results indicate that the optimal design changes as the number of subjects changes. Another set of results confirms what seems to be expected intuitively: as the number of subjects grows, the relatively fixed development cost becomes a smaller and smaller percentage of the total cost. These two sets of results seem to be related directly. Since, for the smaller samples, development costs constitute the majority of total cost, the optimal design includes more raters than prompts. That is, the burden of reliability is shifted to the least expensive (in relative terms) part of the assessment. (Contains 2 figures, 4 tables, and 14 references.) (Author/SLD)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED 400 277

Optimal Designs for Performance Assessments: The Subject Factor

Jay Parkes

The Pennsylvania State University

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

☑ This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

PERMISSION TO REPRODUCE AND DISSEMINATE THIS MATERIAL HAS BEEN GRANTED BY

JAY PARKES

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

Paper presented at the Annual Meeting of the National Council on Measurement in Education: New York, NY. April, 1996.

TM025532

BEST COPY AVAILABLE

Abstract

Much speculation abounds concerning how expensive performance assessments are or are going to be. Recent projections indicate that, in order to achieve an acceptably high generalizability coefficient, many additional tasks may need to be added which will enlarge costs. Such projections are, to some degree, correct and to some degree simplistic. The current investigation uses two synthetic examples, based on published costs and variance components, and a constrained optimization procedure to examine the complex relationships among reliability, cost, and sample size. The results indicate that the optimal design changes as the number of subjects changes. Another set of results confirms what seems to be intuitively expected: as the number of subjects grows, the relatively fixed development cost becomes a smaller and smaller percentage of the total cost. These two sets of results seem to be directly related. Since, for the smaller samples, development costs constitute the majority of total cost, the optimal design includes more raters than prompts. That is, the burden of reliability is shifted to the least expensive (in relative terms) part of the assessment.

Optimal Designs for Performance Assessments: The Subject Factor

Two of the more immediate perceived roadblocks to the implementation of performance assessments are low reliability and high cost. Current understandings of how the two are related, however, may suffer from basically two problems. The first problem is that the two are rarely addressed simultaneously in an empirical fashion. The second problem, which may actually be a cause of the first problem, is the assumption that the two are linked in a Spearman-Brown style relationship. The idea that the relationship is much more complex than that was raised by Sanders, Theunissen and Baas (1989). This investigation attempts to add to the understanding of this complex relationship.

Speculation about cost

Many researchers in the last several years have been finding that performance assessments produce low generalizability coefficients (e.g. Shavelson and Baxter, 1992; Shavelson, Baxter and Gao, 1993; Koretz, Klein, McCaffrey, and Stecher, 1994; Koretz, Stecher, Klein, McCaffrey, and Deibert, 1994; Koretz, Stecher, Klein, and McCaffrey, 1994; McWilliam and Ware, 1994). Furthermore, they have noticed that these coefficients are not due so much to rater variance, which was the scourge of reliability in scoring from the 1960's to the 1980's (c.f. Huot, 1990), but of task variance or task-by-subject variance. Since the g-coefficients were low, some of the researchers (e.g. Shavelson, Baxter and Gao, 1993; McWilliam and Ware, 1994) projected the number of tasks necessary to achieve acceptable (e. g. ≥ 0.80) g-coefficients. These projections were large (as many as 23 science tasks in Shavelson et al., 1993), which then led to the inference that they would be very expensive.

Other researchers have also estimated a large cost. In discussing large-scale portfolio assessment, Reckase (1995) concluded that, compared to current multiple-choice methods, portfolios would be a "very expensive alternative (p. 14)." White (1994) held the opinion that,

although the cost would come in different places (i.e. scoring instead of development), the overall costs would be comparable. Hoover and Bray (1995) to some extent validated this claim by showing that the Iowa Writing Test could be conducted for approximately the same cost as the Iowa Test of Basic Skills, albeit the former covered a much smaller domain than the latter.

Until recently, however, the two problems of low reliability and high cost have been discussed together theoretically but not joined together empirically. When this happens, the relationship is much more complex than it first appears. The assumption that adding more tasks will make the assessment both more reliable and more costly relies on three lines of reasoning which may or may not be appropriate: first, that the relationship between task and reliability is the same as that between number of items and reliability as expressed in the Spearman-Brown Prophecy Formula; second, it is not grounded in an empirical technique which takes both concerns into account simultaneously; and third, it seems to ignore the sample-dependent nature of reliability and cost. In contrast Sanders, Theunissen and Baas (1989) claim that it is actually possible to decrease cost while increasing reliability. They also provided a procedure for optimizing an assessment design, that is, minimizing cost while holding the g-coefficient at or above a given level. Building on that work, Parkes and Suen (1995), using the constrained optimization algorithm of Sanders et al. (1989, 1991, 1992), showed that for any given assessment situation, there are many optimal designs. It is for the designer on site to say which would be optimal given the constraints reasonable to that situation.

The current investigation adds another piece to the understanding of the complex nature of the relationship between cost and reliability. First, the results here indicate that the number of subjects is a situational variable which will alter the optimal design of the assessment. Second, as the number of subjects changes, so do the proportional relationships between development costs, scoring costs, material costs, and the total cost of the assessment.

The Synthetic Assessment Situations

Three different investigations form the basis for the current analyses. The goal was to approximate reliability and cost data based on published studies for two different assessment situations. The first is a limited writing sample situation and the second is a large-scale portfolio assessment.

The cost data for both situations are taken from Hoover and Bray (1995), who report on cost information for an administration of the Iowa Writing Assessment. The assessment tested the writing skills of 30,000 school students from grades three to twelve, each of whom wrote two pieces of writing. Each sample was scored twice holistically and twice analytically. For this assessment, Hoover and Bray estimate that $138,000 was spent in developing the 40 writing prompts; $174,410 was spent to score the prompts; and $30,000 was spent for materials.

In the optimization procedure that is to follow, it is necessary to have an estimate of how much cost each aspect of the situation (rater, subject, prompt) contributes to the total cost. In order to achieve this, base units of development, scoring, and material costs were calculated and then a total cost function constructed. For example, the development cost is dependent on both the number of prompts developed and the number of prompts each subject completes. To obtain a base unit cost for development, the $138,000 development cost was divided by 40 prompts to obtain a development cost of $3450 per prompt, and that was divided by two since each person wrote two. This produces the estimate of $1725 for each prompt that each person has to write. Therefore, the development cost function is $1725n_p$, where n_p is the number of prompts each person must write. The scoring cost ($174, 410) was divided by the number of subjects (30,000), the number of prompts per subject (2), and the number of raters or readings per piece (2) to produce a unit scoring cost of $1.43 per prompt, per rater, per subject. The materials were estimated to cost $1.00 per subject. Therefore, the total cost function is:

$$\text{Total Cost} = \$1725n_p + \$1.43n_pn_rn_s + \$1.00n_s. \qquad (1)$$

Parkes and Suen (1995) produced variance components for fifty subjects writing four prompts which were read by three raters. These variance components, which constitute the first situation, are given in Table 1.

Data from the Vermont Portfolio Project as published in Koretz, Stecher, Klein, McCaffrey & Deibert (1994) were used to create the second situation. The variance components from the Grade 4 writing portfolios were used and are given in Table 2. These components are based on portfolios consisting of two parts read by two raters from a total of 1,714 subjects.

INSERT TABLES 1 AND 2 HERE

The First Situation

The first synthetic example combines the Parkes and Suen (1995) variance components with the Hoover and Bray (1995) cost estimates.

The Variance Model

In the Parkes and Suen (1995) variance model, two facets are fully crossed: writing prompt (p), and rater (r). The object of measurement is student's overall writing ability (s). Thus in the generalizability framework, the variance model is:

$$\sigma^2_{(x_{spr})} = \sigma^2_s + \sigma^2_r + \sigma^2_p + \sigma^2_{sr} + \sigma^2_{ps} + \sigma^2_{pr} + \sigma^2_{psr}. \qquad (2)$$

For the optimization analyses, the relative model of measurement was used. Thus, relative error variances were estimated through:

$$\sigma^2(\delta) = \frac{\sigma^2_{sr}}{n_r} + \frac{\sigma^2_{sp}}{n_p} + \frac{\sigma^2_{srp}}{n_r n_p}, \qquad (3)$$

where n_r and n_p are the number of raters and prompts in each particular scenario respectively. The G-coefficient of interest was thus:

$$E\rho^2 = \frac{\sigma_s^2}{\sigma_s^2 + \sigma^2(\delta)}. \quad (4)$$

The Optimization Procedure

A branch-and-bound integer programming algorithm, which is a linear programming technique, was employed to estimate the optimal combination of raters and prompts. This investigation used the solver function of Microsoft EXCEL, version 5.0, to execute the algorithm. The variance components from Table 1, the cost function given in equation 1, and the number of prompts, raters, and subjects were entered into the EXCEL worksheet.

The following optimization problem was submitted for analysis.

Objective Function: Minimize L = Total Cost = $\$1725n_p + \$1.43n_pn_rn_s + \$1.00n_s$;(5)

Subject to: $$E\rho^2 = \frac{\sigma_s^2}{\sigma_s^2 + \sigma^2(\delta)} \geq 0.8, \quad (6)$$

$$n_p \text{ and } n_r \text{ are integers,} \quad (7)$$

and $$n_p \text{ and } n_r \geq 1. \quad (8)$$

The objective function is to minimize the total cost. Constraint (6) specifies the minimal acceptable level of generalizability. Constraints (7) and (8) further delimit the search to a feasibility region of positive integers.

The results of this analysis produces the optimal number of prompts and raters that will assure minimum cost with a g-coefficient at or above 0.8. This analysis was conducted for sample sizes ranging from 25 to 50,000.

The Second Situation

The second synthetic example combines the Koretz et al. (1994) variance components with the Hoover and Bray (1995) cost estimates.

The Variance Model

In the Koretz et al. (1994) variance model, two facets are used: part (p), and rater (r). The object of measurement is student's overall writing ability (s). Coincidentally, then, the equations for the variance model, the relative error variance and the generalizability coefficient are identical here to those for the Parkes and Suen data. Therefore, the variance model is given in Equation 2; the relative error variance us given in Equation 3; and the generalizability coefficient is given in Equation 4. It is worth noting, however, that a g-coefficient was calculated for each of the five subscales (purpose, organization, details, voice, and mechanics). That is, each subscale has its particular variance components and g-coefficient, as is evident in Table 2.

The Optimization Procedure

As with the first situation, the variance components from Table 2, the cost function given in equation 1, and the number of parts, raters, and subjects were entered into the EXCEL worksheet.

The following optimization problem was submitted for analysis.

Objective Function: Minimize L = Total Cost = $\$1725n_p + \$1.43n_pn_rn_s + \$1.00n_s$;(9)

Subject to: $$E\rho^2 = \frac{\sigma_s^2}{\sigma_s^2 + \sigma^2(\delta)} \geq 0.8, \text{ for each subscale} \qquad (10)$$

$$n_p \text{ and } n_r \text{ are integers,} \quad (11)$$

$$\text{and} \quad n_p \text{ and } n_r \geq 1. \quad (12)$$

The results of this analysis produces the optimal number of prompts and raters that will assure minimum cost with a g-coefficient at or above 0.8. This analysis was also conducted for sample sizes ranging from 25 to 50,000.

Results

In addition to producing an optimal design for each sample size, corresponding development, scoring, material, and total costs were also derived. Tables 3 and 4 contain the optimal designs at selected sample sizes as well as the dollar figures and percentage of total cost attributable to each cost category.

INSERT TABLES 3 AND 4 HERE

For both synthetic examples, the same pattern of results emerges, as is evident in Figures 1 and 2. There is not one optimal answer that holds for all sample sizes. For small samples, the optimal designs contain more ratings than prompts. In contrast, the optimal designs for the larger samples contain more prompts than ratings.

INSERT FIGURES 1 AND 2 HERE

Furthermore, as the sample size increases, the proportions of development, scoring, and material costs to total cost changes. In each case, for smaller sample sizes, development costs

BEST COPY AVAILABLE

represent by far the largest proportion of total cost. But as the samples get larger, scoring costs take on the lion's share of total cost.

The first result mentioned above is influenced by the second result. Since the constrained optimization algorithm is tasked with minimizing cost while maintaining a g-coefficient of 0.8 or better, it will consider relatively lower cost parts in order to boost reliability before considering the higher cost items. That means that when development costs are relatively large, ratings are considered; when scoring costs are relatively large, prompts are considered.

Discussion

This investigation has been designed to shed light on the complexity of the relationship between cost and reliability. In order to do so, synthetic examples combining real data from different sources were used. There are some drawbacks to this approach. Many assumptions had to be made about the cost structures used. Since the variance components and the cost data came from different sources, there is no guarantee that the cost function calculated was the most appropriate to work with the data. In other words, it would be inappropriate to take these optimal designs back to Vermont and implement them. The ideal would be to have cost data from Vermont.

This study is an improvement on mere speculation because it utilizes data from assessments that were actually conducted. It could be improved upon by having both cost and reliability data from the same source. This investigation does gain some generalizability, however, since the results held for two examples, one based on a large-scale assessment and one based on a small-scale assessment.

The focus of this investigation, however, is the processes and relationships involved, not the actual numbers produced. In this regard, some of the results provide some food for thought. The results here seem to indicate that the size of the sample being assessed is an important, if not the most important, factor that determines optimal designs.

Previous speculation that increasing tasks would increase cost appears simplistic in light of these results. By relying on generalizability alone without using any cost data, the picture was clearer: increase tasks. In the optimization framework, both generalizability concerns (such as size of the variance component) and cost concerns are considered. Using this approach, the picture is more complex: check the effect of sample size before turning to tasks.

On the surface, it intuitively sounds counter-productive to add complexity to this issue. At a deeper level, though, the previous simple state of the relationship between cost and reliability had led to an impasse. If increasing tasks was the only way to get a more reliable assessment; and doing so was going to make them even more costly than they already were; performance assessments were between a rock and a hard place. The use of the constrained optimization procedure to simultaneously consider cost and reliability has provided a route through the impasse. It has shown, for example, that many designs can achieve the psychometric constraints. This study has added another piece to this alternate route: when simultaneously considering cost and reliability, the sample size will affect the optimal design. In the present cases, the number of tasks is reduced as sample size increases, thus, at least partially exonerating task as the culprit causing high cost.

References

Hoover, H. D., and Bray, G. (1995, April). *The research and development phase: Can a performance assessment be cost effective?* Paper presented at the Annual Meeting of the American Educational Research Association, San Francisco, CA.

Huot, B. (1990). The literature of direct writing assessment: Major concerns and prevailing trends. *Review of Educational Research, 60 (2)*, 237-263.

Koretz, D., Klein, S., McCaffrey, D., and Stecher, B. (1994). *Interim report: The reliability of the Vermont portfolio scores in the 1992-93 school year.* (RAND/ RP - 260). Santa Monica, CA: RAND.

Koretz, D., Stecher, B., Klein, S., and McCaffrey, D. (1994). The Vermont portfolio assessment program: Findings and implications. *Educational Measurement: Issues and Practice, 13 (3)*, 5-16.

Koretz, D., Stecher, B., Klein, S., McCaffrey, D., and Deibert, E. (1994). *Can portfolios assess student performance and influence instruction?.* (RAND/ RP-259). Santa Monica, CA: RAND.

McWilliam, R. A., and Ware, W. B. (1994). The reliability of observations of young children's engagement: An application of generalizability theory. *Journal of Early Intervention, 18 (1)*, 34-47.

Parkes, J. and Suen, H. K. (1995, April). *The preferability of constrained optimization in determining the number of prompts, modes of discourse, and raters in a direct writing assessment.* Paper presented at the Annual Meeting of the National Council of Measurement in Education, San Francisco, CA.

Reckase, M. D. (1995). Portfolio assessment: A theoretical estimate of score reliability. *Educational Measurement: Issues and Practice, 14(1)*, 12-14.

Sanders, P. F., Theunissen, T. J. J. M., and Baas, S. M. (1989). Minimizing the number of observations: A generalization of the Spearman-Brown formula. *Psychometrika, 54* , 587-598.

Sanders, P. F., Theunissen, T. J. J. M., and Baas, S. M. (1991). Maximizing the coefficient of generalizability under the constraint of limited resources. *Psychometrika, 56 (1)*, 87-96.

Sanders, P. F., Theunissen, T. J. J. M., and Baas, S. M. (1992). The optimization of decision studies. In M. Wilson (Ed.) *Objective Measurement: Theory into Practice (Vol. 1)*. Norwood, NJ: Ablex.

Shavelson, R. J. and Baxter, G. P. (1992). What we've learned about assessing hands-on science. *Educational Leadership, 49 (8)*, 20-25.

Shavelson, R. J., Baxter, G. P., and Gao, X. (1993). Sampling variability of performance assessments. *Journal of Educational Measurement, 30 (3)*, 215-232.

White, E. M. (1994). *Teaching and assessing writing: Recent advances in understanding, evaluating, and improving student performance, 2nd ed.* San Francisco, CA: Jossey-Bass Publishers.

Table 1: Variance Components from Parkes and Suen (1995)

Source	Variance Components
Student	5.778889
St. X Prompt	2.359104
St. X Rater	0.711556
Residual	12.07971
G- Coeff.	0.76

Table 2: Variance Components from Koretz, Stecher, Klein, McCaffrey & Deibert (1994)

Source	Purpose	Organization	Details	Voice	Mechanics
Student	0.147	0.163	0.152	0.158	0.195
Student X Part	0.032	0.040	0.042	0.042	0.056
Student X Rater	0.007	0.006	0.009	0.007	0.006
Residual	0.193	0.204	0.183	0.209	0.172
G-coefficient	0.43	0.43	0.41	0.37	0.46

16

TABLE 3: Parkes and Suen (1995) Data

Subjects	Prompts	Ratings	Development*	Scoring*	Materials*	Total
25	2	26	$3,450 (65)	$1,859 (35)	$25 (0.5)	$5,334
50	3	8	$5,175 (49)	$1,716 (25)	$50 (0.7)	$6,941
100	3	8	$5,175 (59)	$3,432 (39)	$100 (1.1)	$8,707
500	3	8	$5,175 (22)	$17,160 (75)	$500 (2.2)	$22,835
1000	4	5	$6,900 (19)	$28,600 (78)	$1,000 (2.7)	$36,500
1500	4	5	$6,900 (13)	$42,900 (84)	$1,500 (2.9)	$51,300
2000	4	5	$6,900 (10)	$57,200 (87)	$2,000 (3.0)	$66,100
5000	8	2	$13,800 (10)	$114,400 (86)	$5,000 (3.8)	$133,200
10000	8	2	$13,800 (5)	$228,800 (91)	$10,000 (4.0)	$252,600
50000	8	2	$13,800 (1.1)	$1,144,000 (95)	$50,000 (4.2)	$1,207,800

*Gives dollar figures with percentage of total cost in parentheses.

TABLE 4: Koretz, Stecher, Klein, McCaffrey & Deibert (1994) Data

Subjects	Parts	Ratings	Development*	Scoring*	Materials*	Total
25	2	7	$3,450 (87)	$500 (13)	$25 (0.63)	$3,975
50	2	7	$3,450 (77)	$1,001 (22)	$50 (1.1)	$4,501
100	2	7	$3,450 (62)	$2,002 (36)	$100 (1.8)	$5,552
500	2	7	$3,450 (25)	$10,010 (72)	$500 (3.6)	$13,960
1000	3	4	$5,175 (22)	$17,160 (74)	$1,000 (4.2)	$23,335
1500	5	2	$8,625 (27)	$21,450 (68)	$1,500 (4.8)	$31,575
2000	8	1	$13,800 (37)	$22,880 (59)	$2,000 (5.2)	$38,680
5000	8	1	$13,800 (18)	$57,200 (75)	$5,000 (6.6)	$76,000
10000	8	1	$13,800 (10)	$114,400 (83)	$10,000 (7.2)	$138,200
50000	8	1	$13,800 (2.2)	$572,000 (90)	$50,000 (7.9)	$635,800

*Gives dollar figures with percentage of total cost in parentheses.

Koretz et al. Data

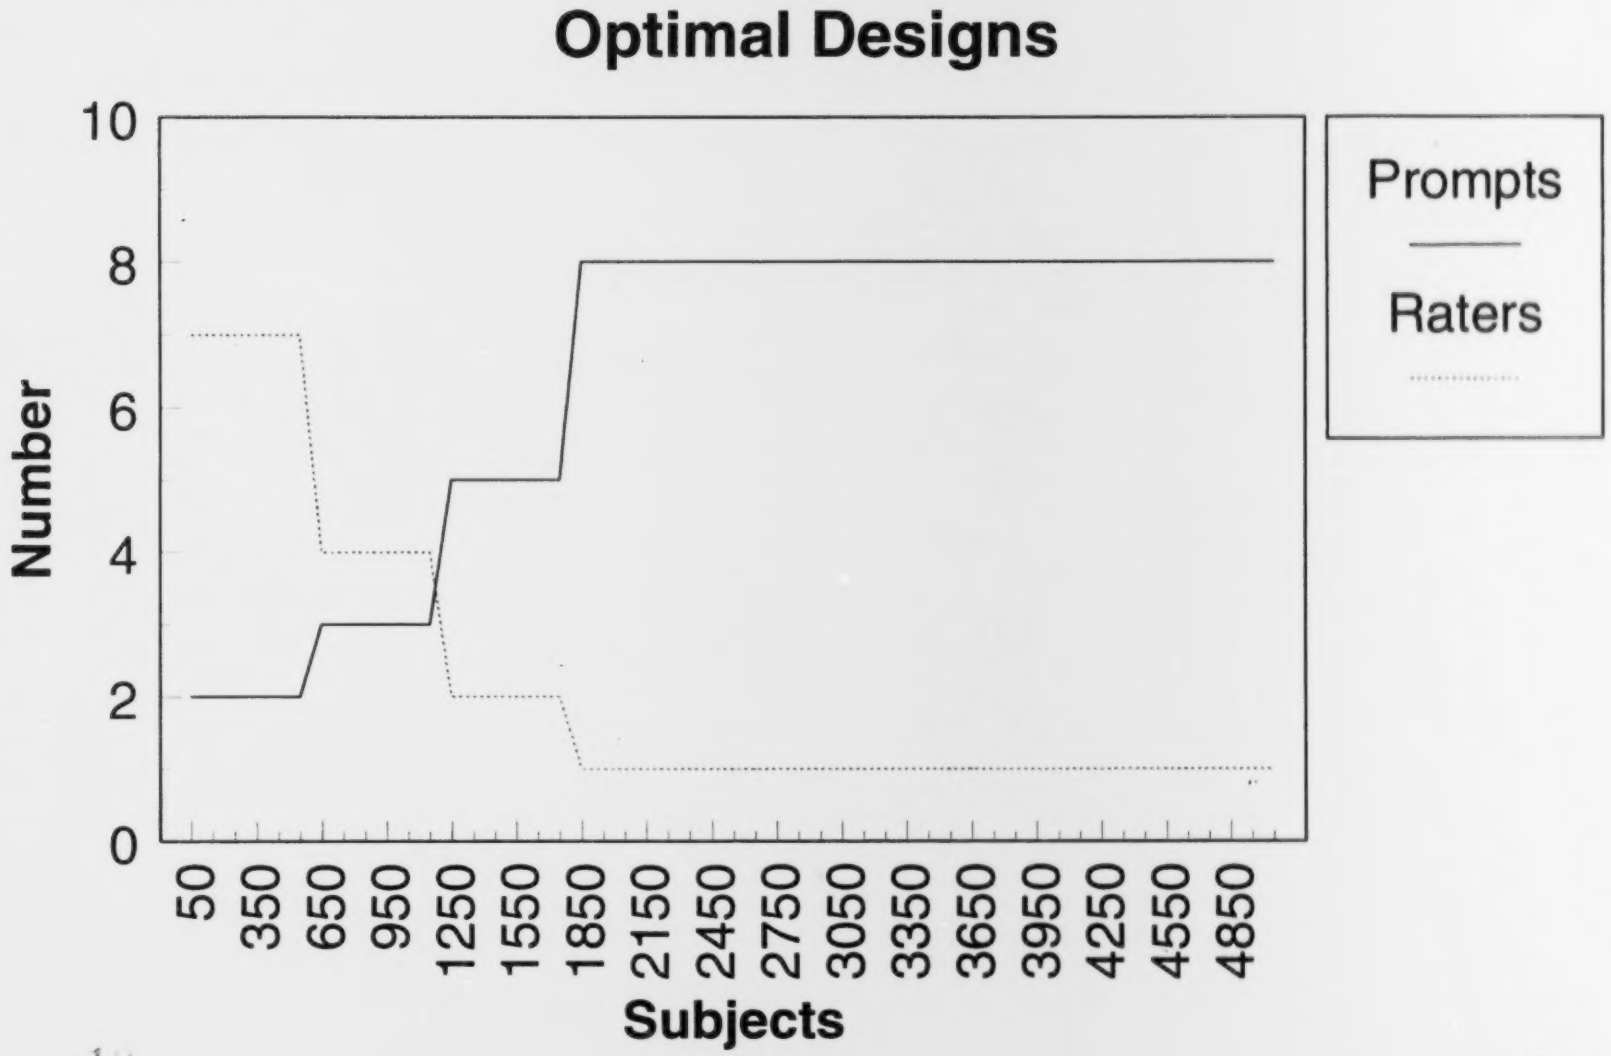

Parkes and Suen Data

Optimal Designs